


IDA DOCUMENT D-1619

AN ALGORITHM FUSION APPROACH TO IRST SIGNAL PROCESSING (III):
NONGAUSSIAN-CORRECTED FILTER FUSION



19950302 119

Jeffrey F. Nicoll
Elizabeth L. Ayers

November 1994

Prepared for
Advanced Research Projects Agency

Approved for public release; distribution unlimited.



INSTITUTE FOR DEFENSE ANALYSES
1801 N. Beauregard Street, Alexandria, Virginia 22311-1772

IDA DOCUMENT D-1619

AN ALGORITHM FUSION APPROACH TO IRST SIGNAL PROCESSING (III):
NONGAUSSIAN-CORRECTED FILTER FUSION

Jeffrey F. Nicoll
Elizabeth L. Ayers

November 1994

Approved for public release; distribution unlimited.



INSTITUTE FOR DEFENSE ANALYSES

Contract DASW01 94 C 0054
ARPA Assignment A-180

Accession For	
NTIS	CRA&I <input checked="" type="checkbox"/>
DTIC	TAB <input type="checkbox"/>
Unannounced <input type="checkbox"/>	
Justification _____	
By _____	
Distribution / _____	
Availability Codes	
Dist	Avail and/or Special
A-1	

PREFACE

The Institute for Defense Analyses was requested to assist the Office of Naval Research (ONR) in the planning and execution of the Infrared Analysis, Modeling, and Measurements Program (IRAMMP). This document summarizes a portion of the work performed with ONR under ARPA Assignment A-180, Infrared Clutter Characterization and Modeling, on alternative algorithms for Infrared Search and Track (IRST) systems from June 1993 to September 1994.

The work was performed in coordination with Mr. Douglas N. Crowder, Naval Surface Warfare Center (NSWC) for the Advanced Research Projects Agency under the technical cognizance of Col. James Silverthorn.

This document has not been subjected to formal IDA review.

CONTENTS

SUMMARY	S-1
I. INTRODUCTION.....	1
II. APPLICATION TO A NONSTATIONARY SCENE	7
III. APPLICATION TO NEARLY STATIONARY SCENES	12
IV. FALSE ALARM LOCATION	18
V. CONCLUSIONS	20
Glossary	GL-1
APPENDIX A—Review of Filter Fusion Approach.....	A-1

FIGURES

1.	IRAMMP Scene 8326-019	3
2.	Signal/Clutter vs. Kurtosis Scatter Plot: Filter 36 (8326-019)	4
3.	Signal/Clutter vs. Kurtosis Scatter Plot: Filter 20 (8326-019)	4
4.	ROC Curves for Filters 20, 36, and 83 (Scene 8326-019).....	8
5.	ROC Curves for Filters 20-36, and 36-83 Fusion (Scene 8326-019).....	8
6.	ROC Curves for Filters 20, 36, and 83 and Their "Linear K" Modifications	9
7.	ROC Curves for 20-36 and 36-83 Fusion, and Their "Linear K" Modifications (Scene 8326-019).....	9
8.	ROC Curves for 20-36 and 36-83 Fusion, and Their "Quadratic K" Modifications (Scene 8326-019).....	10
9.	Four scan Average ROC Curves for Filters 20, 36, and the Optimal Filter and their "Linear K" Modifications (Scene 8326-019).....	11
10.	Four scan Average ROC Curves 20-36 and 36-Optimal Filter Fusion and their "Linear K" Modifications (Scene 8326-019).....	11
11.	IRAMMP Scene 8335-016	12
12.	ROC Curves for Filters 20, 36, and the Optimal Filter (8335-016)	13
13.	ROC Curves for Filters 20, 36, and the Optimal Filter and Their "Linear K" Modifications (Scene 8335-016)	13
14.	Signal/Clutter and Normalized Kurtosis for Targets and False Detections for the Optimal Linear Filter (Scene 8335-016)	14
15.	ROC Curves for Filters 20, 36, and the Optimal Filter and Their "Linear K" Modifications (Scene 8335-016)	14
16.	IRAMMP Scene 8335-006	15
17.	ROC Curves for Filters 20, 36, and the Optimal Filter and Their "Linear K" Modifications (8335-006).....	16
18.	ROC Curves for the Optimal Filter and 36-Optimal Fusion and Their "Linear K" Modifications (8335-016).....	16

19.	Average ROC Curves for Filters 20, 36, and the Optimal Filter With Their "Linear K" Modifications (8335-006)	17
20.	ROC Curves for the Optimal Filter and 36-Optimal Fusion and Their "Linear K" Modifications (8335-006).....	17
21a.	IRAMMP Scene 8335-016. Filter 79 False Detections.....	18
21b.	IRAMMP Scene 8335-016. Filter 79 (Linear K) False Detections	18
22a.	IRAMMP Scene 8326-019. 20-36 Fusion.....	19
22b.	IRAMMP Scene 8326-019. 20-36 Fusion (Linear K)	19

SUMMARY

Clutter is the predominant problem in Infrared Search and Track (IRST) systems. Good range performance in clutter-free environments is easy to achieve at reasonable false alarm rates. However, in the presence of cloud, sea surface, horizon or terrain clutter, there can be a serious increase in false detections with a concomitant decrease in range performance. These clutter environments are particularly troublesome because they do not satisfy the axioms for simple linear signal processing: stationarity and "gaussianity."

The first conventional assumption is that clutter is stationary, that is, that its statistical properties do not vary from one part of the scene to another. This assumption is obviously a gross oversimplification as anyone can see by looking at a cloudy sky. There are obvious changes in character between the clouds and the sky, between one cloud and another, and even within a single cloud. Thus, the assumption of stationarity is often weakened to an assumption of patches of local stationarity separated by nonstationary boundary regions. Since the false detections arise primarily in the nonstationary regions, even a small proportion of nonstationary regions in a scene can cause a serious degradation in performance.

The second assumption is that the statistics of the scene are gaussian. This means, for example, that the distribution of pixel intensities is described by a gaussian (normal) distribution. The assumption of gaussian behavior is applied, in particular, to the results of passing a target detection signal processing filter over the image. These signal residuals would conventionally be considered as a normal distribution of mean zero and standard deviation σ . A pixel is considered a target if the value of the signal (filter output), s , is larger than some threshold, T , times the standard deviation: $s > T\sigma$. If the distribution of residuals were gaussian, then the expected number of false detections could be easily computed from the threshold. However, in real scene, the distribution of residuals is nongaussian, having a long tail that can be described as exponential. Thus, there are more false detections than would be predicted on a gaussian hypothesis.

Both stationarity and gaussianity are necessary for the standard linear matched filter to be optimal—that is, to provide the best probability of detection for a fixed false alarm

rate. The nonstationary and nongaussian character of the clutter makes the problem more difficult and opens the door to alternate techniques that may improve performance.

Another important characteristic of the IRST problem is that it is dominated by low-probability events. IRST systems operate by examining the tails of the distributions and thus are most sensitive to the nongaussian, nonstationary, and nonoptimizable part of the scene. The threshold for the spatial processor in a realistic system might be set to pass only 10^{-5} or 10^{-6} of the available scene pixels. All of these elements of the IRST background contribute to the difficulty of making a system that is robust against changes in the backgrounds. They combine to change the problem from that of a target immersed in a background to a target competing against other discrete objects (edges and linear features as well as point-like) in a background.

This paper is part of a series of papers that attempts to turn the clutter characteristics to advantage.^{1,2} The first approach was to apply concepts used in sensor fusion to the IRST domain. The basic idea is to use two distinct signal processing chains, and compare the results from each. For example, one could use a linear filter in one chain and a non-linear filter in the other. At low thresholds (large numbers of false detections) one would expect the set of false detections found by one method to be essentially the same as those found by the other. However, as the threshold is raised and the total number of detections drops, the overlap between the sets of detections diminishes. By combining the two filters, the number of false detections can be reduced and overall performance improved.

This approach can provide substantial gains when applied to ad hoc filters; the gains would apply to an IRST that had a fixed set of filters. A more stringent test is to compare the filter fusion approach to the standard matched linear filter that is computed for each scene analyzed. Since the scene may have only a limited nonstationary amount, the matched filter, if defined for each image dynamically, can be hard to beat. The filter fusion approach can still provide some advantages over the dynamically generated matched filter, but they are not as dramatic as the gains over a fixed filter set. Filter fusion provides an

- 1 "An Algorithm Fusion Approach to IRST," J.F. Nicoll and E.L. Ayers, *Proc. of the 1992 Targets, Backgrounds and Discrimination IRIS Specialty Group*; "Sensor Fusion with a Single Sensor: An Algorithm Fusion Approach to IRST Signal Processing," J.F. Nicoll and E.L. Ayers, IDA Document D-1156, 1992.
- 2 "Filter Fusion: Using Algorithm Fusion for IRST," J.F. Nicoll, E.L. Ayers, D.N. Crowder, and B. Billard, *Proc. of the 1993 Targets, Backgrounds and Discrimination IRIS Specialty Group*; "Sensor Fusion with a Single Sensor: An Algorithm Fusion Approach to IRST Signal Processing (II)," J.F. Nicoll and E.L. Ayers, IDA Document D-1408, 1993.

improvement in probability of detection at the lowest false alarm rates and an increase in the consistency of performance from image to image.

The second approach to improving the performance over that of the standard filter uses the nongaussian character of the clutter directly. At the 1992 Targets, Backgrounds, and Discrimination Specialty Group Meeting, General Electric (now Martin Marietta) discussed an enhancement of the performance of the F-14D IRST algorithms by making the threshold depend on the local "edginess" as well as the standard deviation of the clutter.³ This showed some performance improvements, perhaps due to the correlation of edges in the scene with the nonstationary character of the clutter near boundaries. In this paper, the same idea was applied to measures of nongaussian behavior rather than edginess, using a threshold that increases in the presence of nongaussian clutter.

Performance was improved for ad hoc filters, the standard matched filter, and for filter fusion. Filter performance became more consistent from scene to scene; if a filter is not well matched to a scene, the nongaussian character of the filter output becomes more prominent and the nongaussian correction to the threshold compensates for the filter's failings. False detections were also reduced by as much as a factor of two at fixed probability of detection. The detections eliminated appear to lie primarily along boundaries and edges—precisely the locations that provide the greatest challenge to the standard approaches. The approach does not always improve performance significantly. However, it never leads to a significant loss in performance.

³ M. Hartless, "Improving the Performance of the F-14 IRST in Severe Clutter Using Adaptive Space-Time Signaling Processing," *Proc. of the 1992 Targets, Backgrounds, and Discrimination IRIS Specialty Group*, Vol. I.

I. INTRODUCTION

Clutter is the predominant problem in Infrared Search and Track (IRST) systems. Good range performance in clutter-free environments is easy to achieve at reasonable false alarm rates. However, in the presence of cloud, sea surface, horizon or terrain clutter, there can be a serious increase in false detections with a concomitant decrease in range performance. These clutter environments are particularly troublesome because they do not satisfy the axioms for simple linear signal processing. The clutter is generally nongaussian and nonstationary; this implies that the matched spatial filter that is the normal approach to signal processing cannot be guaranteed to be optimal. Another important characteristic of the IRST problem is that it is dominated by low-probability events. IRST systems operate by examining the tails of the distributions and thus are most sensitive to the nongaussian, nonstationary, and nonoptimizable part of the scene. The threshold for the spatial processor in a realistic system might be set to pass only 10^{-5} or 10^{-6} of the available scene pixels. All of these elements of the IRST background contribute to the difficulty of making a system that is robust against changes in the backgrounds. They combine to change the problem from that of a target immersed in a background to a target competing against other discrete objects (edges and linear features as well as point-like) in a background.

We recently introduced an approach^{1,2} that attempts to turn the clutter characteristics to advantage. The basic idea is to use two distinct signal processing chains, and compare the results from each. For example, one could use a linear filter in one chain and a non-linear filter in the other. At low thresholds (large numbers of false detections) one would expect the set of false detections found by one method to be essentially the same as those found by the other. However, as the threshold is raised and the total number of detections drops, the overlap between the sets of detections diminishes. Consider, for example, the

- 1 "An Algorithm Fusion Approach to IRST," J.F. Nicoll and E.L. Ayers, *Proc. of the 1992 Targets, Backgrounds and Discrimination IRIS Specialty Group*; "Sensor Fusion with a Single Sensor: An Algorithm Fusion Approach to IRST Signal Processing," J.F. Nicoll and E.L. Ayers, IDA Document D-1156, 1992.
- 2 "Filter Fusion: Using Algorithm Fusion for IRST," J.F. Nicoll, E.L. Ayers, D.N. Crowder, and B. Billard, *Proc. of the 1993 Targets, Backgrounds and Discrimination IRIS Specialty Group*; "Sensor Fusion with a Single Sensor: An Algorithm Fusion Approach to IRST Signal Processing (II)," J.F. Nicoll and E.L. Ayers, IDA Document D-1408, 1993.

case of a single false detection in each processing chain; it is extremely unlikely that the detections will be the same. We therefore propose the following discriminant:

Fusion discriminant. Two signal processing schemes are applied to the same frame of input data. A pixel is declared a detection if the outputs of both signal processing methods pass two separate thresholds.

This method is formally identical to the "AND" method of sensor fusion. Such algorithms work if the two sensors are correlated on the targets and uncorrelated on the false alarms. However, this is very nearly guaranteed for the IRST case:

Target Correlation. The two algorithms are chosen to be "close to optimal" in some sense. This implies that true targets will have strong signatures under both of the schemes that are well localized at the target location.

Clutter Decorrelation. On the other hand, clutter outputs will differ more than the target outputs since the processing chains were designed for the targets not the clutter. The long tails of the clutter distribution imply that the clutter objects in the tail may be sensitive to signal processing changes and the location of a clutter-induced false detection may be unstable under changes in the processing (for example, the detection may "walk" along the edge of a cloud feature as the signal processing is changed).

In the earlier work, this technique was shown to enhance the performance of typical ad hoc linear and nonlinear filters³ and beat the performance of the optimal linear filter, at least at the lowest false alarm rates.⁴ This is possible, of course, only if the scene is nonstationary or nongaussian or both. At the 1992 Targets, Backgrounds, and Discrimination Specialty Group Meeting, General Electric (now Martin Marietta) discussed an enhancement of the performance of the F-14D IRST algorithms by making the threshold depend on the local "edginess" as well as the local variance of the clutter.⁵ This showed some performance improvements, perhaps due to the correlation of edges in the scene with the non-stationary character of the clutter near boundaries.

Since the filter fusion approach depends on the nongaussian character of the filter residuals, it appeared interesting to combine the idea of constant false alarm rate (CFAR) thresholds depending on more than the variance with the filter fusion approach. In this

³ "An Algorithm Fusion Approach to IRST," J. F. Nicoll and E. L. Ayers, *Proc. of the 1992 Targets, Backgrounds, and Discrimination IRIS Specialty Group*.

⁴ "Filter Fusion: Using Algorithm Fusion for IRST," J. F. Nicoll and B. Billard, *Proc. of the 1993 Targets, Backgrounds, and Discrimination IRIS Specialty Group*.

⁵ M. Hartless, "Improving the Performance of the F-14 IRST in Severe Clutter Using Adaptive Space-Time Signaling Processing," *Proc. of the 1992 Targets, Backgrounds, and Discrimination IRIS Specialty Group*, Vol. I.

paper, we consider the use of a threshold that depends on the degree of nongaussianity of the clutter.

Figure 1 shows an MWIR (medium wave infrared) IRAMMP (Infrared Analysis, Modeling, and Measurements Program) cloud scene (8326-019) with a considerable degree of obvious nonstationary character. As a measure of nongaussian behavior, we choose the kurtosis (K), defined as the average of the fourth power of the filter residuals. For simplicity we normalize the computed value of the kurtosis by the expected value for gaussian data. For gaussian data, $K \approx 1$.

$$K = \frac{4 \langle p^4 \rangle}{3\pi^2 \langle |p|^4 \rangle} . \quad (I-1)$$

For each candidate detection the mean absolute value of the filter residual (used as a common surrogate for the standard deviation) and the normalized kurtosis can be calculated in a local region around the target pixel.



Figure 1. IRAMMP Scene 8326-019

Figure 2 shows a scatter plot of the signal-to-clutter ratio (S/C) versus the normalized kurtosis for approximately 500 targets inserted into the image in a widely distributed grid and for the 100 most intense signal/clutter image points for a simple nonlinear (median subtract) filter, denoted as Filter 36. Note that the kurtosis spread of the Filter 36 residuals is enormous compared to that computed in the vicinity of the inserted targets; the kurtosis of clutter points makes them (for this scene and this filter) separable from targets. This depends on the filter applied. Figure 3 shows the scatter plot corresponding to a linear filter (Filter 20) for the same scene. There is still some separation for the Filter 20 residuals, but it is not as marked as for Filter 36 (note the difference in scales). From the scatter plots, one might predict that accounting for the nongaussian component will greatly improve Filter 36 while having a smaller effect on Filter 20.




Figure 2. Signal/Clutter vs. Kurtosis Scatter Plot: Filter 36 (8326-019)




Figure 3. Signal/Clutter vs. Kurtosis Scatter Plot: Filter 20 (8326-019)

This paper will explore the relationship between the nongaussian character and fusion. A systematic study of all the ways to exploit the nongaussian character was not performed. Instead, an examination of the scatter plots suggested several ad hoc threshold manipulation schemes. The first increases the threshold, T , linearly with the normalized kurtosis, K . Thus the required signal-to-clutter ratio to declare a target is given by:

$$\frac{S}{C} > T + K - 1 . \quad (I-2)$$

This will be denoted as a "linear K" approach. Of course, the coefficient of K in Eq. (I-2) could be varied from unity to further improve performance. The value chosen seemed roughly in the right ballpark. A second approach is a quadratic shift in the threshold, referred to as "quadratic K":

$$\frac{S}{C} > T + 2(K - 1)^2 . \quad (I-3)$$

Again, the coefficient was chosen by visual examination of selected scatter plots. The goal was to show that using the nongaussian information could improve performance, not to devise an optimal scheme.

The deviation of the computed kurtosis from the expected gaussian value can be attributed to two different causes. First, the region may be part of a stationary, homogeneous, nongaussian distribution with a deviant value of the kurtosis. Second, and more likely in the current situation, the region may be inhomogeneous. For example, in the vicinity of an edge, the distribution of pixels will be bimodal. The kurtosis is then measuring the separation of the bimodal distribution. If one assumes that the bimodal distribution is represented by two gaussians with standard deviation σ and a separation 2Δ , then the relevant moments are given by:

$$\langle p^4 \rangle = 3\sigma^4 + 6\sigma^2\Delta^2 + \Delta^4 \quad (I-4)$$

$$\begin{aligned} \langle |p| \rangle &= \sqrt{\frac{2}{\pi}} \sigma + \Delta \operatorname{erf}\left(\frac{\Delta}{\sqrt{2}\sigma}\right) - \frac{2\sigma}{\sqrt{2\pi}} [1 - e^{-\Delta^2/2\sigma^2}] \\ &\approx \sqrt{\frac{2}{\pi}} \sigma [1 + \frac{\Delta^2}{2\sigma^2} - \frac{\Delta^4}{24\sigma^4}] \end{aligned} \quad (I-5)$$

where the last line follows for small values of the separation, 2Δ . In this case the normalized kurtosis is less than one:

$$K \approx 1 - \frac{\Delta^4}{\sigma^4} . \quad (I-6)$$

In the limit of large separation, this becomes $K \approx 4/3\pi^2$. For a widely separated unsymmetric bimodal distribution with one mode containing a fraction p of the points,

$$K = \frac{1}{12\pi^2} \frac{p^3 + (1-p)^3}{p^3(1-p)^3} . \quad (I-7)$$

This can give normalized kurtosis values greater than one for $p < 1/5$. Thus the normalized kurtosis, while not an ideal metric for indicating inhomogeneous or bimodal regions, may help in identifying these regions.

II. APPLICATION TO A NONSTATIONARY SCENE

As was shown,⁶ for relatively stationary scenes, the fusion scheme only beats the optimal linear filter in the case of the lowest false alarm rates. However, for a nonstationary scene such as 8326-019, there can be some surprising results. Figure 4 gives the ROC (receiver operating characteristic) curve (probability of detection versus false detections) results⁷ for filters 20, 36, and 83, the optimal filter for the scene. Note that for this scene, the "optimal" filter is in no way optimal and that the nonlinear filter is very poor; the ad hoc linear filter, Filter 20, is fortuitously the best. Figure 5 compares the best filter (Filter 20) with two filter fusions, Filter 20 with Filter 36, and Filter 36 with Filter 83. The fusions are comparable to Filter 20, but are certainly not systematically better. There is a hint of better performance at the lowest false alarm rates (as discussed in Hartless, 1992).

Figure 6 compares Filters 20, 36, and 83 with their "linear K" versions. Note that there is a slight improvement in Filter 20 and the optimal filter but an enormous improvement in Filter 36, which is now also better than the optimal filter. Figure 7 compares the fusions of the "linear K" algorithms with the best previous filter fusion results. The linear K fusions are clearly superior at all the false detection rates shown to both the previous fusions and the optimal linear filter. Figure 8 shows that comparable results are obtained if the quadratic K modification is made, with the linear K approach being slightly better.

The results are fairly encouraging. A scene with a considerable degree of non-stationarity is partly tamed by the use of the nongaussian extensions to the filters and is further improved by filter fusion. The result shows considerable improvement over the previous best results for the scene and clearly surpasses the optimal filter defined for the entire scene. Of course, a scene segmenter that used different filters in different domains might provide more competitive performance, but the margin gains from filter fusion and

⁶ Hartless, 1992.

⁷ Note that if these results are compared to Hartless, 1992, the target strength has been lowered 20% in this work. The shapes of the ROC curves and the comparisons are not strongly dependent on the target strength; stronger targets just lift the ROC curves upward.




Figure 4. ROC Curves for Filters 20, 36, and 83 (Scene 8326-019)




Figure 5. ROC Curves for Filters 20-36, and 36-83 Fusion (Scene 8326-019)




Figure 6. ROC Curves for Filters 20, 36, and 83 and Their "Linear K" Modifications (Scene 8326-019)




Figure 7. ROC Curves for 20-36 and 36-83 Fusion, and Their "Linear K" Modifications (Scene 8326-019)




Figure 8. ROC Curves for 20-36 and 36-83 Fusion, and Their "Quadratic K" Modifications (Scene 8326-019)

the nongaussian threshold adjustment might still apply, especially in the boundary regions between segmented regions. These are the primary source of false detections and are not simply handled by a simple segmentation scheme.

These results hold relatively robustly for all the scans of this scene. Filter 20 is usually the best filter with the optimal filter coming closer to it in performance in other scans. Filter 36 is abysmal in all the scans examined. Figure 9 shows an average ROC curve over 4 different scans of scene 8326-019 comparing the performance of Filters 20, 36, and the optimal filter computed for each scan with their linear K modifications. Again the modification has only a small improvement for the linear filters but makes a marked improvement in the nonlinear filter, Filter 36. Figure 10 compares the average ROC curves for the fusion results. The best results are achieved by fusing the filters with nongaussian threshold corrections. Comparing the figures shows Filter 20 (the best average single filter) lies between the fusion results of the ordinary and linear K filters. Thus both fusion and the nongaussian threshold adjustment are required to get the maximum performance.




Figure 9. Four scan Average ROC Curves for Filters 20, 36, and the Optimal Filter and Their "Linear K" Modifications (Scene 8326-019)




Figure 10. Four scan Average ROC Curves 20-36 and 36-Optimal Filter Fusion and Their "Linear K" Modifications (Scene 8326-019)

III. APPLICATION TO NEARLY STATIONARY SCENES

In this second, the method is applied to scenes that are more nearly stationary and gaussian. These are scenes for which the optimal linear filter defined for the whole scene is a good performer and is hard to beat with the ad hoc filter fusion approach. For such scenes the nongaussian component of the filter residuals is generally smaller and the advantages of using a kurtosis-adjusted threshold scheme might be expected to be less. Figure 11 shows an IRAMMP scene, 8335-016, which shows less structure than the scene used previously (Fig. 1).




Figure 11. IRAMMP Scene 8335-016

The stationarity or near stationarity of this scene is reflected in the ROC curves of the filters. Figure 12 shows the results for Filter 20, Filter 36, and the optimal filter for this scene, Filter 79. In this case, Filter 79 is clearly the best with Filter 36 slightly better than Filter 20.

As shown in Fig. 13, the use of the kurtosis-adjusted threshold has little effect on Filters 20 and 36, but markedly improves the optimal filter (Filter 79) giving it a flatter ROC curve and lifting the performance at the lower false detection rates. Figure 14 shows the scatter plot of signal/clutter and normalized kurtosis for the Filter 79 residuals. The improvement does not require a large separation of the target and false detection scatter distributions.

Fusing the filters does not lead to any further improvement. For this scene the kurtosis-adjusted optimal filter seems to be the best. However, some of the fusions are close in performance, such as the 36-79 linear K fusion, which is better for all false detection rates greater than 8; this finding apparently persists for all higher rates.




Figure 12. ROC Curves for Filters 20, 36, and the Optimal Filter (8335-016)




Figure 13. ROC Curves for Filters 20, 36, and the Optimal Filter and Their "Linear K" Modifications (Scene 8335-016)




Figure 14. Signal/Clutter and Normalized Kurtosis for Targets and False Detections for the Optimal Linear Filter (Scene 8335-016)




Figure 15. ROC Curves for Filters 20, 36, and the Optimal Filter and Their "Linear K" Modifications (Scene 8335-016)

Figure 16 shows another relatively unstructured, nearly stationary scene, taken about 90 minutes earlier on the same day as the preceding scene.




Figure 16. IRAMMP Scene 8335-006

As shown in Fig. 17, the optimal filter (Filter 77) is clearly superior, with the advantage in performance being the smallest at low false detection rates. In this case, Filter 20 appears to be generally superior to Filter 36.

Applying the kurtosis adjustments does not markedly change the behavior of the single filters except that the optimal filter's performance at low false detection rates is again improved. The probability of detection rises from less than 0.4 to greater than 0.5 at 1 false detection per scan. The fusion results are shown in Fig. 18 and show no real change in performance. These equivocal results are relatively consistent in this scene. Figures 19 and 20 show the average performance of the single and fused filters along with their analogous kurtosis-adjusted filters for 4 neighboring scans of the same scene. The addition of the kurtosis to the threshold produces a marginal improvement. The fusion of the nonlinear filter (36) and the optimal filter shows some advantage at the lowest false detection rate but then has weaker performance. The lowest detection rate performance may be the more significant and small benefits may persist and increase as the detection rate is lowered.




Figure 17. ROC Curves for Filters 20, 36, and the Optimal Filter and Their "Linear K" Modifications (8335-006)




Figure 18. ROC Curves for the Optimal Filter and 36-Optimal Fusion and Their "Linear K" Modifications (8335-016)




Figure 19. Average ROC Curves for Filters 20, 36, and the Optimal Filter With Their "Linear K" Modifications (8335-006)




Figure 20. ROC Curves for the Optimal Filter and 36-Optimal Fusion and Their "Linear K" Modifications (8335-006)

IV. FALSE ALARM LOCATION

The use of the kurtosis can reduce the false alarm rate by about a factor of two at fixed probability of detection. Figure 21a and Fig. 21b show the locations of the false alarms with and without the kurtosis correction for scene 8335-1662 and the optimal filter. For about the same probability of detection (set at 0.671) there are 5 detections with the kurtosis correction and 11 without. The false detections removed by the kurtosis correction lie primarily along cloud edges and boundaries. The detections clustered around the dark hole on the right side of the image are removed. Note that two of the false detections in Fig. 21a and one in Fig. 21b lie along a bad detector line.




Figure 21a. IRAMMP Scene 8335-016. Filter 79 False Detections




Figure 21b. IRAMMP Scene 8335-016. Filter 79 (Linear K) False Detections

Figures 22a and 22b show the comparison for Scene 8326-1904 between 20-36 fusion with and 20-36 fusion without the linear kurtosis correction. The use of the kurtosis correction removes 6 of the 12 false detections, again primarily along the

boundaries of cloud structures. Note that the cluster of detections around the "hole" has been largely removed and the remaining detections appear uncorrelated with any obvious structure.




Figure 22a. IRAMMP Scene 8326-019. 20-36 Fusion




Figure 22b. IRAMMP Scene 8326-019. 20-36 Fusion (Linear K).

This suggests that the kurtosis correction is sensitive to the boundaries between inhomogeneous regions in the background image. This may be important if the kurtosis correction approach is applied to more sophisticated signal processing. For example, one method of dealing with inhomogeneous background clutter is to automatically segment the image into more nearly homogeneous regions and apply a locally defined optimal linear filter to each region. Such an approach should have better performance than a more globally defined "optimal" filter. However, the boundaries between segmented regions will still be difficult to treat since they do not fall naturally into a single filter domain. The use of the kurtosis correction may improve the performance of the signal processing by reducing the false detections along the boundaries of each region.

V. CONCLUSIONS

We have expanded the work done previously on the filter fusion approach to IRST by explicitly considering nongaussian effects. The conclusions can be summarized as follows:

1. False detections for a variety of scenes and filters may have a separable different distribution of nongaussian characteristics than for the scene as a whole. This provides a possible discriminant for reduction of false alarms.
2. A simple measure of nongaussian behavior, namely the kurtosis, can be used to reduce false detections by raising the threshold in the vicinity of high values of the kurtosis. The performance improvements do not seem to depend on the precise analytic forms of the threshold-kurtosis relation.
3. Other measures were also considered (but not discussed in this paper); for example, neither the skewness nor the absolute value of the skewness of the distribution of filter residuals offered any advantages. The normalized 6th power of the filter residuals provided performance improvements essentially identical to the kurtosis.
4. The improvement in performance is greatest for inhomogeneous scenes. In this case the nongaussian measure may be responding to nonstationarity rather than local nongaussian distributions. Initial efforts to isolate a more direct measure of the nonstationarity have not been as satisfactory as the kurtosis.
5. False detections may be reduced by as much as a factor of 2 at fixed probability of detection. The detections removed appear to lie primarily along boundaries. This indicates that if a more elaborate segmentation approach is used to define homogeneous regions, the kurtosis correction could have value along the boundary of such regions. It may be that a measure based directly on bimodality would perform as well or better than the kurtosis since the kurtosis may not be as sensitive as a direct measure of inhomogeneity. Initial work on a direct measure of bimodality shows some promise.
6. The kurtosis corrections can improve the performance both of filter fusion approaches and standard optimal linear filters.
7. The approach does not always improve performance significantly. However, it does not ever lead to a significant loss in performance.

Further work will include a comparison of this approach with approaches using bimodality or edge strength to mark the regions of most probable false detections in an effort to improve the computational efficiency and performance of the approach.

GLOSSARY

ARPA	Advanced Research Projects Agency
CFAR	constant false alarm rate
IFOV	instantaneous field of view
IRAMMP	Infrared Analysis, Modeling, and Measurements Program
IRST	Infrared Search and Track
NSWC	Naval Surface Warfare Center
ONR	Office of Naval Research
ROC	receiver operating characteristic

APPENDIX A

REVIEW OF FILTER FUSION APPROACH

APPENDIX A

REVIEW OF FILTER FUSION APPROACH

This appendix reviews the basic filter fusion approach and summarizes the salient results from the earlier work.

All of the algorithms applied here to cloud data take the form of 1-dimensional spatial filters that combine a high-pass filter to suppress low-frequency clutter and a target filter matched to the expected form of the target signature. An estimate of the background clutter level is made using the average absolute value of the filter residuals and a threshold is applied to the signal-to-clutter ratio. A variety of combinations of filter elements were explored. The ones discussed in this document are given in the following table. The common target filter is (0.15, 0.5, 0.85, 1.0, 0.85, 0.5, 0.15).

Table A-1. Filter Descriptions

Filter No.	Mean/Clutter	Background Estimator
20	MR2	B2
36	Anti-Median (31)	B2

MR2 denotes a simple one-dimensional second difference filter linear filter, modified to reflect the over-sampling of the IRAMMP sensor.

MR2 = (-1/8, -1/8, -1/8, -1/8, 0, 0, 1, 0, 0, 1/8, -1/8, -1/8, -1/8).

B2 is 3 x 75 with a 23-sample inner window on the center channel and one sample wide on the adjacent channels for a total of 200 samples in the estimate.

Pixels exceeding this threshold ("pixel exceedances") are merged to form false detections.¹ The fusion algorithm requires the selection of two thresholds for the two different signal processing chains. In any sensor fusion scheme, this choice of the relationship between the thresholds is always a subtle one. Here an order statistics approach to the problem of threshold choice is used. This has the advantage of being robust when the

¹ Merging of adjacent pixel exceedances to form a single detection is particularly important in IRAMMP data, which is sampled at approximately 3.25 to 1 in the scan direction.

signal processors have different distributions of residuals² and requires no prior knowledge of the correlation between the processing approaches. The top N detections for each process are compared; any detection that appears on both lists is a detection of the fused system for the pair of thresholds determined by the threshold of the Nth detection for process 1 and the Nth detection for process 2.

In Table A-2, we compare two signal-processing chains defined by their filters; there are no pixels in common until $N = 15$, two by $N = 24$, five by $N = 34$, etc. The total number of false detections of the fused process when the first 100 false detections of each individual process are compared is 23, a reduction by a factor of more than four.

Note that the clutter estimate used to define the threshold is given by the average absolute value and is not converted to a pseudo-gaussian standard deviation. This is contrary to the usual convention but since the filter residuals are not gaussian, any such conversion does not have any real significance. When comparing these values of the threshold with other literature, however, the conversion must be made; the conversion factor being ≈ 0.8 , a threshold of 5 in this work corresponds to a threshold of approximately 4 in "sigma" units.

The filter residuals, both before and after fusion, are well described by an exponential distribution; the fusion of the two filters reduces the amplitude and increases the slope by a factor of approximately 1.67. This slope increase is expected in AND fusion: for the ideal case of completely uncorrelated processes, the slopes would increase by a factor of 2. This increase in slope is important because it shows that the benefits of fusion increase as the threshold increases, that is, as the false alarm rate is decreased. The reduction of the thresholds is not a complete description of the sensor fusion process. A true target may pass one of the two thresholds and not the other: there is a loss due to the AND fusion and therefore a ROC curve analysis must be made; the ROC curve gives the probability of detection versus the false detection rate and is the true measure of the algorithm's performance.

² If both distributions were gaussian and mean-zero, it might be appropriate to take equal thresholds as the definition of the fused algorithm. However, even this method does not guarantee the optimum results. See the discussion in "Annual Summary Report for FY 1990, Task T-D2-748, Assessment of Advanced Sensor Systems, Volume I: Issues in Automatic Target Recognition," IDA Document D-923, J.F. Nicoll, J.D. Silk, and D.A. Sparrow, 1991.

Table A-2. Fusion Thresholds for Scene 8335-006

Fused Detections	N	Threshold for Filter 20	Threshold for Filter 36
1	15	5.20	5.50
2	24	5.07	5.15
5	34	4.95	5.03
10	51	4.80	4.85
15	65	4.69	4.74
20	78	4.61	4.64
23	96	4.52	4.53

Figure A-1 shows the number of false detections as a function of threshold value for Filters 20 and 36 and for 20-36 fusion. Note that at each threshold the number of detections is reduced and that the slope of the fusion thresholds is approximately 1.7 times greater than for the individual filters. For an ideal fusion case (zero correlations between the false detections of the two filters), the increase in the slope would be a factor of 2.

Figure A-2 shows the ROC curves for Filters 20 and 36 and for 20-36 fusion. Note that the probability of detection decreases only very gradually as the false detection rate is lowered for the fusion case.




Figure A-1. Threshold Relations for Filters 20 and 36 and their Fusions




Figure A-2. ROC Curves (Probability of Detection vs. False Detections) for Filter 20, Filter 36 and 20-36 Fusion

Fusion improves the performance of ad hoc filters. A stronger test is to compare the fusion approach with the "best" single filter approach. Figure A-3 shows the ROC curves for Filter 77, the conventional optimal linear filter (an optimized version of Filter 20), and Filter 78, which uses the antimedian filter used in Filter 36 as a preprocessor and then applies the optimal linear filter computed from the residuals (an optimized version of Filter 36). The optimal linear filter is the best single filter; the nonstationary aspects of this scene are not severe and therefore the optimal linear filter is a very good choice for a filter. Note also that the optimized antimedian filter (Filter 78) is better than the plain antimedian filter (Filter 36).

Figure A-4 shows the optimal linear filter compared with the fusions of Filter 20 with Filter 36, and Filter 36 with Filter 77. For false detection rate of more than three false detections per image, the optimal linear filter is still the best. However, for the smaller detection rates, the fusion algorithm is better than the optimal filter! It is at the smaller detection rates that the comparison should be made. An IRAMMP image has 128 rows and 1452 columns with an over-sampling rate in the scan direction of about 3.25 to 1. Thus, there are approximately $5.7 \cdot 10^4$ independent detection opportunities. A false detection rate of 10^{-5} corresponds to less than one false detection per IRAMMP image. Therefore, the interesting part of the ROC curve is the low false detection end.




Figure A-3. ROC Curves (Probability of Detection vs. False Detections) for Filter 20, Filter 36, Filter 77, and Filter 78 applied to Scene 8335-006




Figure A-4. ROC Curves (Probability of Detection vs. False Detections) for Filter 77, 20-36 Fusion, and 36-77 Fusion applied to Scene 8335-006

The other possible fusions are not shown but are not as good as the ones displayed. The optimized antimedian filter, although better than the plain antimedian filter when considered as a single filter, is a worse fusion partner; similarly, the fusion of the two optimized filters (77 and 78) is not as good as the fusions shown. It may be that the linear optimization increases the correlation of the false detections between the filter candidates.

These results suggest that the value of fusion increases as the false detection rate is lowered. Four successive scans in the same direction were selected from IRAMMP scene 8335-006 and analyzed as if they formed a single image of 4 times the size (scans 16, 18, 20, and 22 were used; the original single frame used in the previous calculations was scan 28). Figure A-5 shows the individual ROC curves for this combined image where Filter 104 and 105 denote the optimal linear filter and optimized antimedian filter (in this case computed from the second of the four scans). Because four scans are used, one can effectively go to a false detection rate of 1/4 false detection per IRAMMP image; Filter 36 is better than Filter 20 and the optimal linear filter is the best single filter.




Figure A-5. ROC Curves for Filter 20, 36, 104, and 105 Applied to Four Scans From Scene 8335-006

Figure A-6 compares the optimum linear filter to the 20-36 and the 36-104 fusions. For this combined scene all of these filters are essentially indistinguishable for false detection rates of 1-5 false detections per IRAMMP image, but the fusions are superior at

the lowest false detection rates. These results show that the fusion scheme using generic ad hoc filters such as Filter 20 and Filter 36 is better than or as good as the optimal linear filter computed for the scene at the lower (and relevant false detection rates) but the optimal linear filter will usually become better than the fusions at sufficiently high false alarm rates.




Figure A-6. ROC Curves for Filter 104 and 20-36 and 36-104 Fusion Applied to Four Scans From Scene 8335-006

One of the factors involved in combining several images is the scan-to-scan variability of the performance of each filter. Figure A-7 shows the performance of Filter 20 on a large number of scans from 8335-006 (scans 15 through 30 omitting 24 and 26 that were marred by severe artifacts). For false detection rates less than 5 detections per image, there is considerable variability between the performance forming a broad ROC "banana." The P_d at 1 false detection varies between 0.3 and 0.8 with scans 16, 20, and 29 being the worst and scans 21, 25, and 27 being the best. This variability is expected since one is dealing with the tails of the distributions. The P_d at 1 false detection is governed by the largest signal-to-clutter value of the clutter; the distribution of the largest value in any collection of 10^5 points is not well behaved! Note that since this is actual sensor data the outlier points are always "real" but may be associated with actual clutter features or sensor anomalies. The results are filter dependent; for Filter 36, the range in P_d is similar but the bad scans are 15, 19, and 29 and the good scans are 16, 21, and 23.




Figure A-7. ROC Curves for Filter 20 (14 scans from Scene 8335-006)

Figure A-8 shows the results of fusing Filters 20 and 36 on the same scans. The ROC banana is tighter with the P_d at one false detection ranging from 0.55 to 0.85. The good scans are 18, 22 and 28 and the bad scans are 15, 19 and 30. Since the fusion method uses more of the tail of the distribution (recall that the first fused detection often corresponds to the 10th detection of the single filters) it is not surprising that fusion approach stabilizes the performance.

All of the calculations presented above have been for single images or multiple images treated as a single unit. The range of all possible temporal processing or tracking algorithms is beyond the scope of this paper. A "stack and add" approach was applied to the set of four scans from IRAMMP scene 8335-006 used previously. This is a primitive temporal processing algorithm that presumes a target motion and simply aligns a sequence of images to that motion and adds the images. The target is summed coherently while the noise adds incoherently; the coherency of the clutter depends on the clutter temporal and spatial correlation properties and the velocity of the target. Since four scans are used, the target strength for the inserted targets was lowered by a factor of 2 to give roughly the same signal-to-clutter ratios.




Figure A-8. ROC Curves for 20-36 Fusion (14 scans from Scene 8335-006)

For the purposes of demonstration the targets were displaced by 5 samples horizontally between each of the images used (since the background was drifting with respect to the sensor, the relative target-clutter velocity was approximately 5 samples per image in the vertical direction). The instantaneous field of view (IFOV) of the IRAMMP sensor is about 0.22 mr with a 3.25 over-sampling rate in the scan direction. Therefore, this displacement corresponds to approximately 0.32 mr radians. Using every other scan gives a mean separation time of 0.564 seconds or an angular velocity of 0.59 mr/sec. This is equivalent to a transverse velocity of 30 m/sec at a range of 50 kilometers. For a target with a speed of approximately 300 m/sec this corresponds to an angle of 5 degrees off the nose; for 100-km range, this would correspond to 10 degrees off the nose. Figure A-9 shows the comparison for the single filters for this stack-and-add example; Filters 110 and 111 are the optimum linear filter and the optimized antimedian filter for the stacked scene. In this case, the antimedian filters outperform the optimal linear filter slightly and Filter 20 is again distinctly inferior. Figure A-10 shows the comparison between Filter 110 and 111 and the fusion of Filter 36 with 110 and Filter 110 with 111. In this case, the 110-111 fusion is superior, providing some advantage over Filter 111 alone (notably remaining flat while 111 has begun to drop) and a considerable advantage over the linear optimal filter.




Figure A-9. ROC Curves for Filters 20, 36, 110 and 111: Four Stack-and-Add Scans (16-18-20-22) from Scene 8335-006: 5 Samples/Scan Target Velocity




Figure A-10. ROC Curves for Filters 36, 110 and 36-111 Fusion: Four Stack-and-Add Scans (16-18-20-22) from Scene 8335-006: 5 Samples/Scan Target Velocity

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public Reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)			2. REPORT DATE November 1994	3. REPORT TYPE AND DATES COVERED Final—June—September 1994
4. TITLE AND SUBTITLE An Algorithm Fusion Approach to IRST Signal Processing (III): Nongaussian-Corrected Filter Fusion			5. FUNDING NUMBERS C - DASW01 94 C 0054 ARPA Assignment A-180	
6. AUTHOR(S) J. Nicoll, E. Ayers				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Institute for Defense Analyses 1801 N. Beauregard St. Alexandria, VA 22311-1772			8. PERFORMING ORGANIZATION REPORT NUMBER IDA Document D-1619	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Advanced Research Projects Agency 3701 N. Fairfax Drive Arlington, VA 22203-1714			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 180 words) <p>The reduction of background clutter in Infrared Search and Track (IRST) systems remains a challenging problem. False alarm inducing clutter is typically nonstationary and nongaussian and may not be adequately treated by linear matched filter techniques. Previously, we examined the simultaneous use of two filters, "filter fusion", to reduce the false alarms. This method depends upon the assumption that the residuals of two sufficiently different filters will be uncorrelated. In the present work we make explicit use of the nongaussian character of the clutter, using a threshold that depends on the local nongaussian character. The introduction of a nongaussian correction often improves the performance of both linear and nonlinear filters, including the optimal linear filter constructed for the scene. Fusing the nongaussian adjusted filters gives a further increase in performance. The use of the nongaussian correction appears to increase the robustness of the filters and the filter fusion, i.e., the performance from frame to frame within a given scene and from scene to scene is consistently improved and made more regular.</p>				
14. SUBJECT TERMS infrared search and track, clutter, signal processing			15. NUMBER OF PAGES 44	16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT SAR	